arXiv:1408.6269v2 [math.NA] 16 Jul 2015

Exploiting Active Subspaces to Quantify Uncertainty in the Numerical Simulation of the HyShot II Scramjet

P.G. Constantine[a,*], M. Emory[b,b], J. Larsson[c,c], G. Iaccarino[b,b]

[a] *Department of Applied Mathematics and Statistics, 1500 Illinois Street, Colorado School of Mines, Golden, Colorado, 80401*
[b] *Center for Turbulence Research, 488 Escondido Mall, Building 500, Stanford University, Stanford, California, 94305*
[c] *Department of Mechanical Engineering, 3149 Glenn L. Martin Hall, University of Maryland, College Park, MD, 20742*

Abstract

We present a computational analysis of the reactive flow in a hypersonic scramjet engine with focus on effects of uncertainties in the operating conditions. We employ a novel methodology based on *active subspaces* to characterize the effects of the input uncertainty on the scramjet performance. The active subspace identifies one-dimensional structure in the map from simulation inputs to quantity of interest that allows us to reparameterize the operating conditions; instead of seven physical parameters, we can use a single derived active variable. This dimension reduction enables otherwise infeasible uncertainty quantification, considering the simulation cost of roughly 9500 CPU-hours per run. For two values of the fuel injection rate, we use a total of 68 simulations to (i) identify the parameters that contribute the most to the variation in the output quantity of interest, (ii) estimate upper and lower bounds on the quantity of interest, (iii) classify sets of operating conditions as safe or unsafe corresponding to a threshold on the output quantity of interest, and (iv) estimate a cumulative distribution function for the quantity of interest.

Keywords: uncertainty quantification, active subspace, hypersonic, scramjet

1. Introduction

Over the past decade there has been a renewed interest in numerical simulations of high-speed air-breathing propulsion systems for hypersonic vehicles, driven by fundamental advancements in simulation tools for high Mach-number flight conditions. These tools are essential in the design process of such vehicles due

*Corresponding author
Email addresses: paul.constantine@mines.edu (P.G. Constantine), memory@stanford.edu (M. Emory), jola@umd.edu (J. Larsson), jops@stanford.edu (G. Iaccarino)
URL: http://inside.mines.edu/~pconstan (P.G. Constantine)

to the challenges and costs associated with traditional physical prototyping. In light of the limited experience with sustained hypersonic flight, numerical simulations can be of critical value to understanding the behavior of this system, in particular characterizing the safe operability limits of the propulsion system.

Supersonic combustion engines (scramjets) are an economic alternative to rockets because they do not require on-board storage of the oxidizer. The HyShot II scramjet, shown in Figure 1, was designed to demonstrate supersonic combustion during flight in a simple configuration [24, 53], and it has since been the subject of multiple ground-based experimental campaigns in the High Enthalpy shock tunnel Göttingen (HEG) of the German Aerospace Center (DLR) [18, 19, 22, 23, 33–36, 51]. The simplicity of the configuration and the availability of experimental data has made the HyShot II scramjet the subject of multiple computational investigations, based on either Reynolds-averaged Navier-Stokes (RANS) simulations [46] or large-eddy simulations (LES) [5, 17, 32].

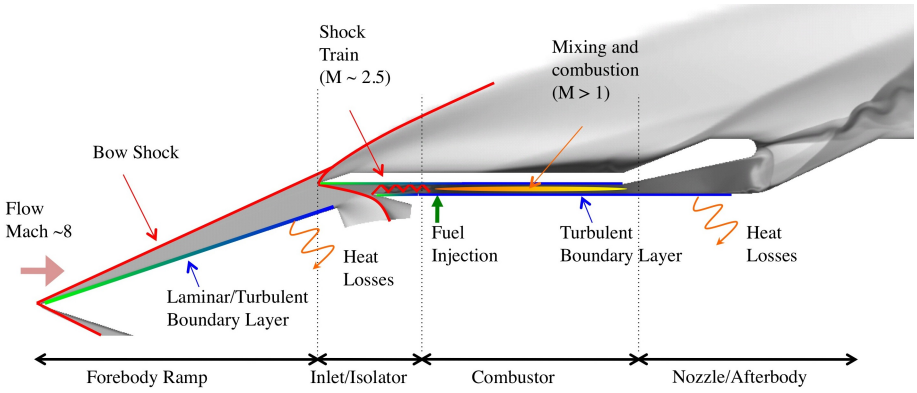

Figure 1: Side view of the HyShot II scramjet outlining the various physics within the different parts of the geometry. Contours show numerical Schlieren results from a 2D simulation.

Although the geometry depicted in Figure 1 (and in detail in Figure 2) is simple and without moving parts, accurately predicting the internal flow structure is a very challenging problem. The difficulty results from the physical phenomena encountered within the engine when traveling at high velocity—namely turbulence, shocks, boundary layers, mixing, and combustion—all of which must be modeled. Required closures and disparate spatial and temporal scales of the various physical processes cause these simulations to be fairly expensive.

Scramjet operation requires a careful balance between maximizing thrust and maintaining stable and safe operation. At the conditions of interest here, the flow inside the HyShot II scramjet combustor has three different regimes [32, 34]. At low fuel-air equivalence ratios (about $\lesssim 0.38$), the flow is supersonic throughout. At higher fuel-air equivalence ratios, a stable flow with a shock-train in the combustor develops. Finally, at sufficiently high fuel-air equivalence

ratios, the flow will unstart—a potentially catastrophic failure mode associated with a large loss in thrust. Scramjet design and operation is complicated by the fact that the maximum thrust occurs very close to the boundaries between these different regimes.

Previous efforts [42, 58, 59] studied various aspects of this problem both experimentally and computationally, but they largely ignored the effect of the uncertainties on the estimation of the operability limit. These efforts enhance understanding of the physics but do not provide sufficient confidence in the quantitative estimates of the maximum fuel flow rate compatible with safe operations. As a consequence safety margins based on experience are applied to the estimate. To produce more trustworthy uncertainty estimates, it is necessary to explicitly identify, estimate, and account for uncertainties that affect the operability limit. Many sources of uncertainty are present in the scenario of interest; both variability in the operating environment of the vehicle (e.g., fluctuations in temperature) and inadequacy of the constitutive models (e.g., the combustion model) introduce uncertainty in the simulation result. Our prior work studies the effects of uncertainties in the RANS turbulence models and how to account for them in a physically meaningful manner [15]. We have also compared the effects of uncertainties in the combustion chemistry to effects of uncertain operating conditions for HyShot II, and we found that the effects of operating conditions were much greater than the effects of the chemistry parameters [10]. The current paper focuses on the HyShot II operating environment vehicle studied by DLR in the HEG shock tunnel [18, 19, 51]. In this context, the main source of variability corresponds to incomplete knowledge and controllability of the conditions within the HEG shock tunnel. A more comprehensive study that includes the effects of the operating condition variability and the uncertainties introduced by the physical models is currently ongoing and will be presented in a forthcoming paper.

Monte Carlo and related random sampling methods are commonly used to quantify uncertainty in a simulation's predictions—e.g., with moments, quantiles, or empirical density functions—given uncertainty in its inputs [44, 54]. Estimates are unbiased, and one can develop confidence intervals for the uncertainty analysis from the Central Limit Theorem. However, when the simulation is expensive, sampling can be impractical for accurate statistics due to the slow $\mathcal{O}(n^{-1/2})$ convergence rate (where n is the number of samples). An alternative approach for rapidly converging, biased estimation employs response surfaces—where one uses a few carefully selected runs of the expensive model to construct a cheaper response surface that is subsequently sampled. Popular response surface models in uncertainty quantification include polynomial chaos [21, 65], stochastic collocation [64], and Gaussian process regression [31, 48]. For polynomial variants, moments reduce to numerical quadrature on the integral formulation of averages. Response surfaces are most appropriate when the number of input parameters is sufficiently small and the outputs are sufficiently smooth functions of the inputs.

We supplement the quantity of interest's statistical measures with its range; determining the range is posed as optimization. We employ *active subspaces* [8]

to discover that the quantity of interest can be well represented by a univariate function of the *active variable* derived from the model's inputs. We then exploit this low-dimensional approximation to (i) estimate the cumulative distribution function of the quantity of interest, (ii) determine the range of the quantity of interest, and (iii) identify the sets of parameters corresponding to safe operation of the scramjet. As a byproduct, the active subspace reveals a subset of the model's parameters whose perturbations produce the greatest change in the quantity of interest, i.e., a sensitivity analysis. For this particular scramjet model, discovering the low-dimensional structure and quantifying uncertainty requires enough simulation runs to fit and validate a global, least-squares-fit linear approximation of the quantity of interest as a function of scramjet inputs. In practice, the cost scales linearly with the dimension of the input space—much lower than the scaling for more complicated approximation procedures. This cost is remarkably small considering both the simulation's complexity and its dependence on seven independent input parameters.

The remainder of this paper is broadly structured in four sections: describing the HyShot II simulation in Section 2, characterizing the uncertainties in the operating conditions in Section 3, outlining the technique for discovering the active subspace in Section 4, and applying the active subspace to quantify uncertainties in Section 5. We conclude with a brief summary.

2. Methodology for HyShot II simulations

2.1. Geometry and computational grids

Figure 2 shows the HyShot II geometry including the intake ramp (forebody), isolator, combustion chamber, and nozzle details. The actual HEG system is much larger and includes the mounting structure and the instrumentation of the $1:1$ scale model. The geometry in Figure 2 corresponds to one-half the flight vehicle, which had symmetrically mounted fueled and un-fueled engines on top of a rocket. Both dual and single engine models have been investigated in the HEG [18, 19, 51].

The intake ramp is significantly wider than the combustion chamber, and thus the flow entering the isolator/combustor is very close to uniform in the spanwise direction. The gap between the intake ramp and body wall of the combustion chamber is a boundary-layer and isolator shock (emanating from the cowl wall leading edge) bleed channel. Within the combustor there are four fuel injection ports spanning the width of the channel, shown in the slice view `A-A`. The model was mounted at a nominal angle-of-attack of $3.6°$.

In order to reduce the computational cost of each simulation the HyShot II configuration is separated into two domains. The first is a 2D representation of the intake ramp and entrance to the combustion chamber, denoted as forebody and isolator in Figure 1, respectively. The second domain is a fully 3D representation of the combustion chamber and exit nozzle, including the fuel injection ports. This domain decomposition assumes that the combustion chamber inflow conditions, i.e., the flow dynamics upstream of fuel injection,

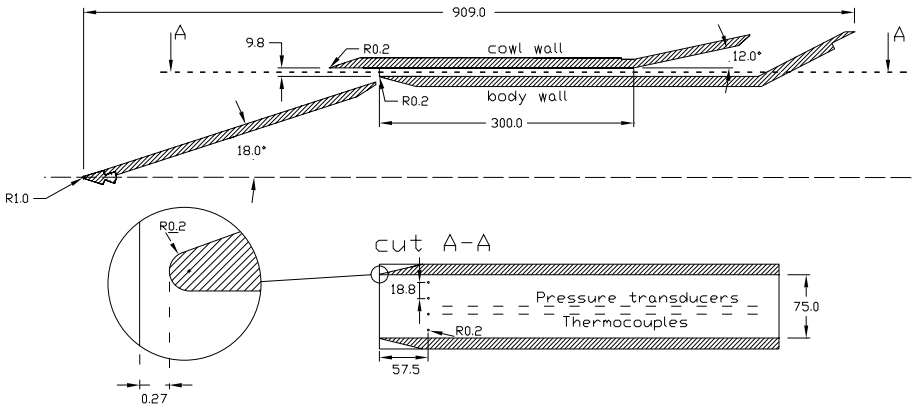

Figure 2: HyShot II scramjet geometry, taken from [46].

are two-dimensional (no variation in the spanwise dimension). This assertion has been verified by DLR through comparison of 2D and 3D intake ramp simulations [29]. Leveraging these observations, a similar domain decomposition has been applied by a variety of researchers numerically investigating the HyShot II scramjet [cf. 17, 29, 30].

The domain decomposition allows us to address the different modeling requirements in the 2D and 3D domains. Over the forebody and through the isolator, turbulence and transition phenomena dictate the flow structures. In contrast, the combustion chamber requires modeling of mixing and flow structures due to the injected fuel. The 2D simulations are dramatically less expensive due to relaxed mesh resolution and modeling complexity. We next describe the meshes used in this analysis.

2.1.1. 2D forebody

The grid for the 2D forebody/ramp is shown in Figure 3, where the downstream boundary is located at the fuel injection ports. In the wall-normal direction the grid is designed such that the first row of cells adjacent to the wall have $y_1^+ = 1$, resulting in just under 50k control volumes in the domain. In regions far from solid boundaries the cells are unstructured, which reduces numerical artifacts related to poor shock-grid alignment.

2.1.2. 3D combustion chamber

To reduce the cost of simulating the 3D domain we take advantage of symmetries and simulate only $1/8^{th}$ of the total combustor. In the work of Pečnik *et al.* [45, 46] the results using a 1/2–span domain (including the sidewall and two injection ports) and a 1/8–span domain (including half an injection port and using symmetry planes in the spanwise directions) were compared. Those authors concluded that while shocks emanating from the sidewalls influence certain quantities of interest, in general the 1/8–span domain is acceptable for use

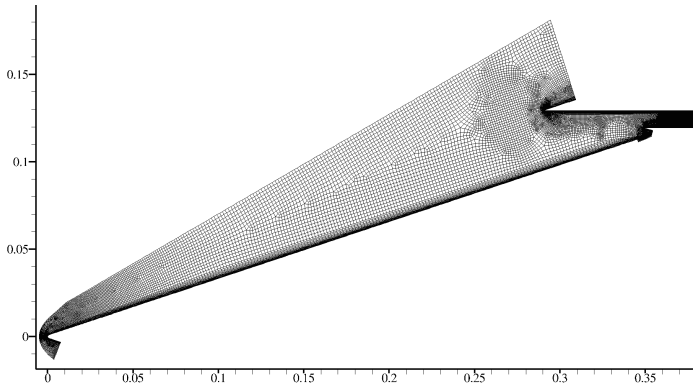

Figure 3: The 2D forebody and isolator mesh used for the HyShot II simulations.

in the analysis of the scramjet. The same conclusion has been reached by DLR and other analysts, all of whom perform HyShot II simulations using a 1/8–span domain [cf. 6, 17, 29, 30].

The grid is again designed such that the wall-normal spacing adjacent to the wall is $y_1^+ = 1$, and the total number of control volumes in this domain is 1.2M. To increase numerical stability, a corner radius of 0.06mm is used at the fuel nozzle orifice; see Figure 4(a). The domain is essentially structured except near the injection port; see Figure 4(c).

2.1.3. Coupling the two domains

In the 3D domain the inflow condition for the oxidizer stream is taken from the 2D simulation. A wall-normal profile is extracted at $x = 352.68$mm (absolute coordinates of the mesh), this profile is applied uniformly across the span of the combustion chamber inlet. The profile is extracted at a location where the oblique shock generated by the body wall leading edge is captured well above the boundary-layer resolving cells, see Figure 5.

2.2. Reynolds-averaged Navier-Stokes

Simulating the HyShot II scramjet with large-eddy simulation (LES) is tractable but costly [5, 32]. Wall-modeled LES uses approximately 100M cells, and wall-resolved LES would require 100B cells. It is not feasible to use LES for uncertainty quantification, which requires many simulations to estimate statistics of the solution behavior. Reynolds-averaged Navier-Stokes (RANS) costs less than LES by modeling only the time-averaged flow. Also, the mesh can be much coarser than LES, which further reduces the cost. The trade-off is that turbulence is now modeled; the accuracy of the result depends strongly on the accuracy of the turbulence model. RANS represents the most practical approach for performing UQ of the HyShot II system, and in general RANS is the most common approach used by hypersonic vehicle designers for simulating turbulence in their systems [29, 66].

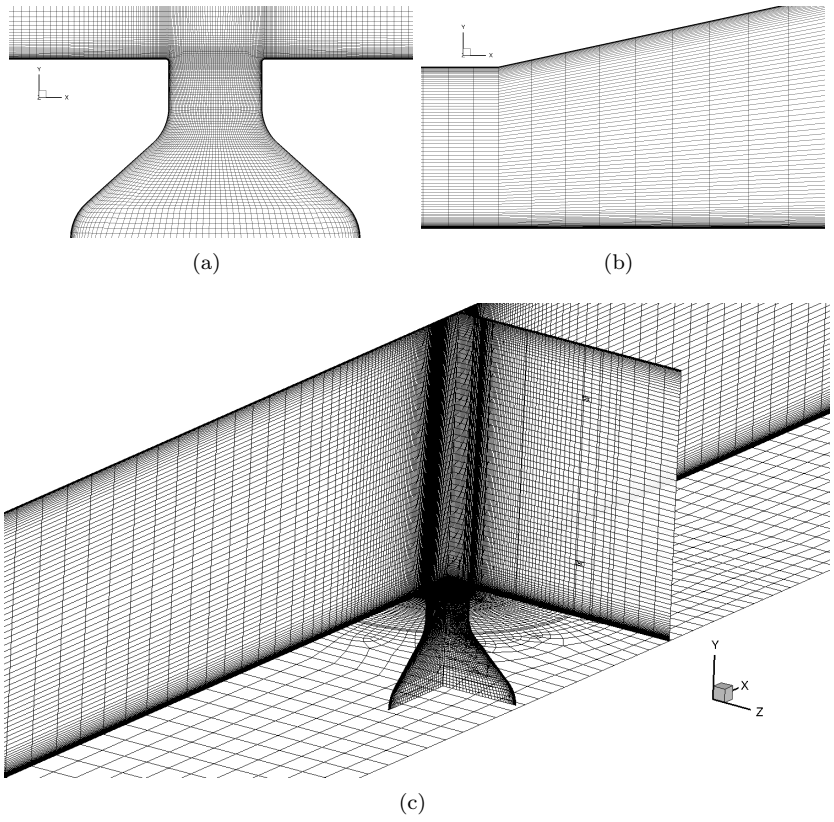

(a) (b) (c)

Figure 4: The 3D combustion chamber mesh used for the HyShot II simulations, highlighting (a,c) the injection port and (b) the nozzle.

Pečnik et al. [46] and Terrapon et al. [55] give comprehensive treatments of the discretization and numerical implementation of the corresponding physical models. Quantitative comparisons between the computational predictions of the pressure within the HyShot II combustion chamber are reported in Pečnik et al. [46]; these comparisons show remarkable agreement in terms of the compression ratio due to fuel injection and burning. Moreover, shock wave locations and strength within the chamber compare favorably. Overall, the simulations capture the dominating effects of the turbulent boundary layer interacting with the shocks. Furthermore, the results compare wall heat flux with experiments, again demonstrating satisfactory agreement. Pečnik et al. [46] also present detailed grid resolution analysis, and they study of the effects of spanwise domain size and side walls; the solver and the results presented in [46] have been used to design the simulations presented in this paper.

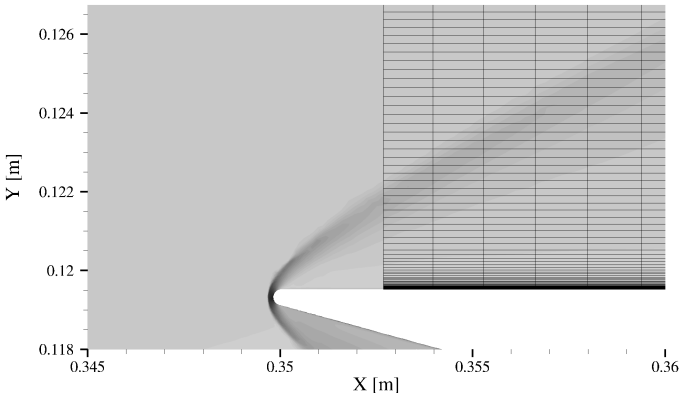

Figure 5: Density contours from the 2D domain overlaid by the 3D mesh (black lines) showing where the 3D inflow profile is extracted relative to the shock structures.

2.2.1. Physics modeling and closures

Several RANS closures are needed in the multiphysics model. We next describe our specific choices.

Turbulence. The turbulence model used to determine the Reynolds stresses is based on the eddy-viscosity hypothesis and the k–ω shear stress transport (SST) formulation [38]. The SST model is one of the most frequently used RANS models in industrial applications. In many applications with compressible boundary layers and shock-turbulence interactions, SST gives reasonably accurate predictions [15, 43]. In the SST model, two transport equations are solved in addition to the RANS equations to describe the turbulent kinetic energy k and the specific dissipation rate ω (having units of inverse time). The SST model blends the standard $k-\omega$ and $k-\epsilon$ models to accurately represent both the near wall regions and the response to high strain and pressure gradient.

Due to its popularity, several modified SST models have been proposed. We apply two specific limiter modifications [47]. The first is a limiter for the eddy-viscosity [39],

$$\mu_t = \frac{\rho a_1 k}{\max\left(a_1\omega, \Omega F_2\right)}, \tag{1}$$

where ρ is the density, $a_1 = 0.31$ is a model constant, k is the turbulence kinetic energy, ω is the specific dissipation, and Ω is the vorticity magnitude. F_2 is a function designed to be 1 for boundary-layer flows and 0 for shear layers. The second is a turbulence kinetic energy production limiter, which is meant to ensure realizable Reynolds stresses,

$$\mathcal{P}_k = \min\left(\mathcal{P}_k, 20\, C_\mu\, k\, \omega\right) , \tag{2}$$

where $\mathcal{P}_k$ is the turbulence kinetic energy production, $C_\mu = 0.09$ is a model constant, and k and ω as above.

Transition. The laminar-turbulent boundary-layer transition process is modeled by manually specifying transition locations on the intake ramp, body wall, and cowl wall. These locations then affect the flow through inhibited production and destruction of turbulence kinetic energy in the boundary layers upstream of these points; this mimics the handling of transition in the $\gamma - Re_{\theta t}$ transition model by Menter et al. [40].

Combustion. The combustion model is based on a flamelet/progress variable approach (FPVA), in which the chemistry is tabulated as a series of laminar flamelet solutions for a given set of boundary conditions and background pressure. The effect of turbulence on the flame is approximated by using a presumed beta probability density function for the fuel/air mixture fraction. In this manner the chemical composition is mapped *a priori* with respect to a small number of parameters used to search this table. The major assumption behind this approach is that chemistry is fast relative to the mixing time scales and can therefore be accurately represented by a small number of scalar quantities. This approach requires three additional transport equations for the mixture fraction, mixture fraction variance, and progress variable. These are the values used with the table to provide species mass fractions and other properties that in turn influence the local temperature and pressure [45, 46, 55]. The 20 reaction $H2-O2$ mechanism of Hong et al. [26] is used to generate the FPVA table. The FPVA framework is fairly common in subsonic application, as it has been developed based on a low Mach assumption. Formal descriptions of the development and extension to high speed flows are found in Saghafian et al. [49], Terrapon et al. [56].

Boundary conditions. In both the 2D and 3D domains, solid walls are modeled as isothermal with $T_w = 300\text{K}$ due to the very short test time which prevents the steel walls from heating up [19]. The nozzle (combustor exit), bleed channel, and freestream (flow outside cowl side of vehicle) use a Neumann boundary condition. At the fuel inlet, the stagnation pressure $P_{0,\text{H2}}$ and the stagnation temperature $T_{0,\text{H2}}$ are specified. The latter is always 300K, whereas the former is varied depending on the desired fuel equivalence ratio.

2.2.2. Flow solver

The present calculations solve the steady, compressible RANS equations (five PDEs), the k–ω SST turbulence model (two PDEs), and the FPVA combustion model (three PDEs and a look-up table). The computations are carried out under the steady-state assumption although the unstart process is transient. This choice is motivated by the desire to detect the conditions that lead to unstart rather than model the entire unstart process. The flow solver `Joe`, developed at Stanford's Center for Turbulence Research, is used to perform these simulations. The code performs parallel calculations on a collocated unstructured mesh using a finite volume formulation. The discretization is second order in space; gradients of flow quantities are calculated with least-squares. Comprehensive discussions of the numerical implementation as well as validation of `Joe`

for HyShot II simulations can be found in Pečnik et al. [45, 46], Terrapon et al. [55].

Cost. The average time for each 2D forebody simulation (Section 2.1.1) used in the uncertainty quantification study was approximately 8.1 hours on 48 CPUs. The average time for each 3D combustor simulation (Section 2.1.2) was approximately 76 hours on 120 CPUs, though there was significant variability in the number of iterations needed to meet the desired stopping criteria.

3. Uncertainty sources and the quantity of interest

The first step in uncertainty quantification is to identify and mathematically describe the system's uncertainties. These uncertainties must be propagated through the model to assess their impact on the quantity of interest. In this section we introduce and characterize both the sources of uncertainty and the quantity of interest in the HyShot II model. We first identify reasonable ranges for the input parameters. All available knowledge (e.g., from observations, theory, and expert opinion) is used to inform the input ranges. These ranges are sufficient to estimate (i) minimum and maximum values for the output quantity of interest and (ii) safe sets of operating conditions. To estimate a complete cumulative distribution function on the output quantity of interest, it is sufficient to consider a density function on the input space. We follow Jaynes' maximum entropy principle [27] to choose a uniform density on the space of inputs bounded by the ranges. Without any additional information (e.g., from actual measurements), the maximally ignorant or least informed assumption is that of a uniform density, which maximizes the (mathematical) entropy over all possible densities on the hyperrectangle defined by the ranges. The uniform density is a modeling choice, and the uncertainty quantification results depend on this choice. A complete study should analyze the results' sensitivity to the choice of density, but this is outside the scope of the present work. The parameters and associated ranges are summarized in Table 2.

3.1. Inflow conditions, mean quantities

The objective is to model conditions from the ground test experimental campaign carried out by researchers at DLR in the HEG. The conditions in the HEG represent the flight test from Boyce et al. [3] along with boundary conditions and comprehensive measurements in the combustor. Stanford University and DLR collaborated to share experimental results, simulation results, and additional unpublished information. When this study was initialized, a total of thirteen experimental runs in the HEG shock tunnel at the relevant conditions (i.e., mimicking the high altitude conditions encountered during the flight test) had been performed, four without fueling and nine with fueling. For each experimental run, the stagnation pressure P_0, temperature T_0, and enthalpy H_0 of the nozzle supply region are measured (reported in Table 1). The fueled runs

	Shot	P_0 [bar]	T_0 [K]	H_0 [MJ/kg]	P_{H2} [bar]	ϕ
fuel-off	805	178.05	2742	3.25		
	807	181.19	2777	3.30		
	808	175.27	2716	3.21		
	814	176.64	2735	3.24		
fuel-on	804	172.96	2652	3.22	5.41	0.341
	809	177.80	2753	3.27	5.27	0.329
	810	176.66	2705	3.20	5.73	0.351
	811	178.81	2726	3.23	4.68	0.286
	812	179.84	2729	3.23	5.32	0.325
	816	173.06	2769	3.39	3.71	0.266
	817	176.99	2796	3.13	5.28	0.315
	827	170.23	2701	3.19	5.09	0.324
	828	187.43	2796	3.28	4.84	0.288

Table 1: Experimental conditions from HyShot II ground tests performed in the HEG [20].

additionally measured the hydrogen fuel plenum pressure $P_{0,\mathrm{H2}}$, from which the fuel/air equivalence ratio ϕ was estimated as

$$\phi = \frac{8\dot{m}_{\mathrm{H2}}}{\dot{m}_{\mathrm{O2}}} \,. \tag{3}$$

The equivalence ratio is a natural quantity to consider given its direct chemical interpretation, but in this study we will consider the fuel plenum pressure as the control parameter since this was the directly measured quantity in the experiments. For all runs the hydrogen plenum temperature was 300K.

Table 1 reports the measured and estimated parameters from the 13 experimental runs. It is important to point out that the stagnation temperature T_0 and stagnation enthalpy H_0 are dependent parameters, being related by the specific heat c_p. When plotting these quantities in Figure 6, it is clear that 4 of the data points are contaminated by errors from some unknown source. The remaining data points (from 9 runs) fall closely on a straight line, which is consistent with the expectation of an approximately constant c_p under these conditions. The least-squares regression model for these 9 data points is

$$H_0 = \frac{T_0 - 508.1386}{6.8718 \times 10^{-4}} \,. \tag{4}$$

We thus take the P_0 and H_0 values from Table 1, and use (4) to calculate the associated T_0. We transform these stagnation conditions into the static pressure, temperature, and velocity to use as inflow conditions in the simulation. This is

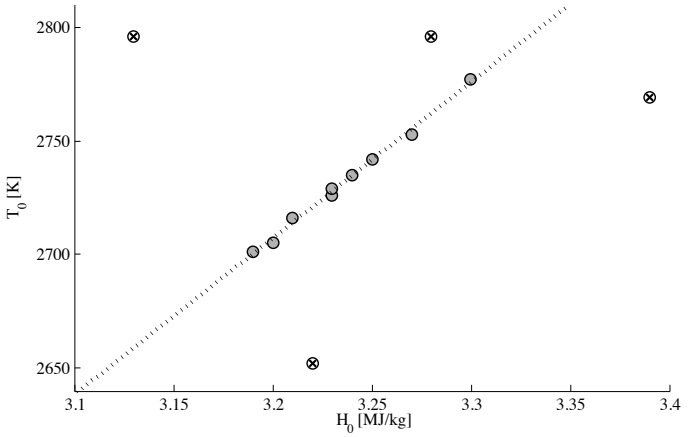

Figure 6: Least-squares fitting (dotted line) of T_0 vs. H_0 data from Table 1, where the experiments marked with $\otimes$ (runs 804, 816, 817, and 828) have been excluded from the curve fit.

done using the following ratios provided by DLR

$$\frac{P}{P_0} = 1.16 \times 10^{-4}\,, \tag{5a}$$

$$\frac{T}{T_0} = 0.0978\,, \tag{5b}$$

$$\frac{U_{\text{mag}}}{\sqrt{H_0}} = 1.332\,, \tag{5c}$$

which are assumed valid for any run condition. Equation (5c) gives the velocity magnitude, which together with the angle-of-attack α allows us to compute velocity components.

While the simulation takes the velocity components and the static pressure and temperature to specify the inflow condition, it is more consistent to characterize the uncertainty directly in the stagnation conditions (P_0 and H_0) that are measured in the experiments. The small number of measurement data available (the 13 runs reported in Table 1) makes it difficult to characterize this uncertainty. Nevertheless, 13 samples in the table are used to compute sample means ($\widehat{P}_0$ and $\widehat{H}_0$) and standard deviations ($\widehat{\sigma}_{P0}$ and $\widehat{\sigma}_{H0}$). The uncertainty range is conservatively defined as $\widehat{P}_0 \pm 3 \cdot \widehat{\sigma}_{P0}$ and similarly for H_0.

We also take the angle-of-attack α as uncertain, due to both the physical mounting of the model in the tunnel and the flow-induced deviation from this angle during a run. The nominal angle-of-attack was 3.6°. Based on the expert opinion of an experienced experimentalist [M. Gamba, personal communication, July 5^{th}, 2012], we believe α can only be specified within $\pm 1°$. Additional evidence suggests this value is uncertain based on analysis performed during the construction of the HyShot II model [18]. The weight of the model and

associated instrumentation caused static load deflection of the model up to 0.465°, a value which will most likely increase during dynamic loading. Also, the entire building shakes while HEG is operating, which further supports the choice model α as uncertain. One may reasonably assume that the model or test section is at an effective angle-of-attack relative to the flow (we ignore any possible uncertainty in the yaw angle of the vehicle). Thus the uncertainty range is specified as $3.6 \pm 1°$. Table 2 summarizes the nominal value and endpoints of the uncertainty range for these parameters.

3.2. Inflow conditions, turbulence quantities

The k–ω SST model requires specification of the turbulence kinetic energy k and the specific dissipation-rate ω at the inflow. This is done by specifying a turbulence intensity I and turbulence dissipation length scale $L_{t,\omega}$ (characteristic of the energy-carrying eddies). Thus,

$$k = \frac{3}{2}(U_{\text{mag}}\, I)^2\,, \tag{6a}$$

$$\omega = \frac{\sqrt{k}}{C_\mu^{1/4} L_{t,\omega}}\,, \tag{6b}$$

where U_{mag} is the mean velocity magnitude and $C_\mu = 0.09$ [57].

There are no direct measurements the turbulence intensity and length scale—neither their nominal values nor their deviation from nominal. We therefore estimate both the nominal values and the uncertainty. Assumed constant stagnation pressure and pressure fluctuation data—specifically that $P'_{\text{rms}}/P_0 = 10^{-5}$—are used to estimate the nominal turbulence intensity as $I = 1\%$. Combining this with expert opinion [M. Gamba, personal communication, July 5^{th}, 2012], we arrive at a range of $I = [0.1, 1.9]$.

The estimates of turbulence length scale are informed by expert opinion [M. Gamba and J. Larsson, personal communication, July 5^{th}, 2012]. The largest turbulence length scales generated at the throat preceding the nozzle in the tunnel are roughly half the size of the throat. Furthermore, 3D isotropic eddies should grow as the inverse of the density in the isentropic expansion in the nozzle,

$$\frac{L_t}{L_{t,\text{throat}}} \approx \left(\frac{P_0}{P}\frac{T}{T_0}\right)^{\frac{1}{3}}. \tag{7}$$

The flow in the nozzle should approximately satisfy the 1D variable-area flow relation

$$\left(\frac{A}{A^*}\right)^2 = \frac{1}{M^2}\left[\frac{2}{\gamma+1}\left(1+\frac{\gamma-1}{2}M^2\right)\right]^{(\gamma+1)/(\gamma-1)}, \tag{8}$$

where γ is the ratio of specific heats. To achieve Mach 7.4 (the nominal HEG flow condition) the nozzle must have an area ratio of $\frac{A}{A^*} = 133$. The diameter of the test section at the HyShot II model is approximately 610mm, thus the throat diameter is 53mm. Substituting the pressure and temperature ratios from (5)

into (7), we find a length scale ratio of 9.43, which leads to a nominal length scale $L_{t,\omega 0} \approx 245$mm. A conservative uncertainty range is specified relative to this nominal value where $L_{t,\omega} = L_{t,\omega 0} \pm 50\%$. Table 2 summarizes the nominal value and endpoints of the uncertainty range for these parameters.

3.3. Transition locations

There are three transition locations within the HyShot II: along the intake ramp/forebody (location labeled $x_{t,\mathrm{r}}$), along the body wall ($x_{t,\mathrm{b}}$), and along the cowl wall ($x_{t,\mathrm{c}}$). Along the forebody $x_{t,\mathrm{r}}$ is determined from heat flux measurements at the vehicle surface. The resolution of thermocouples is very coarse, however, with only four thermocouples spanning 115mm [20]. The body and cowl transition locations are specified through expert opinion (based on the experience of the HEG facility operators) and there are no transition measurements available on these walls. Based on this information, the three transition locations are assumed uncertain—though are still treated as fixed locations (ignoring uncertainty related to the assumption that these locations are stationary).

In this work we consider uncertainty in $x_{t,\mathrm{r}}$ and $x_{t,\mathrm{c}}$, both of which reside in the 2D computational domain. Note that $x_{t,\mathrm{r}}$ is defined relative to the vehicle leading edge and $x_{t,\mathrm{c}}$ is defined relative to the cowl leading edge; these locations are shown in Figure 3. The 2D simulations are significantly cheaper than their 3D counterpart, so the impact of considering additional uncertainties is relatively small. The 3D computations are more expensive due to both increased domain size and the additional modeling closures (e.g., combustion, mixing). Thus uncertainty in $x_{t,\mathrm{b}}$ is ignored as a practical consideration to reduce computational cost of the uncertainty quantification.

DLR uses a simple transition model to estimate the location $x_{t,\mathrm{r}}$=145mm because of the limited resolution provided by the thermocouples mounted on the forebody walls; this estimate is consistent with experimental correlations that indicate transition occurs where

$$\frac{Re_\theta}{M_e} \approx 200\,, \tag{9}$$

where Re_θ is the Reynolds number based on momentum thickness and M_e is the Mach number at the edge of the boundary-layer. As opposed to directly assuming that the magnitude of $x_{t,\mathrm{r}}$ is uncertain, we instead assume that the critical value of $Re_\theta/M_e = 200$ from transition theory is uncertain.

It is well known (see, e.g., [40]) that the transition Reynolds number is sensitive to several factors including the free-stream turbulence, the wall roughness, etc. It is therefore difficult to specify a precise threshold value. In our previous uncertainty analysis [47], we consider the effect of compressibility on the growth of turbulent spots in the transition region by introducing uncertainty in the correlation specification in (9). Here we follow the same strategy and define that transition is equally likely to occur at the point where

$$\frac{Re_\theta}{M_e} = 200\,(1 \pm \varphi), \tag{10}$$

where $\varphi = 0.2$. The ratio Re_θ/M_e is linear in the momentum thickness θ, which grows as $\theta(x) \propto \sqrt{x}$ in the laminar boundary layer prior to transition. Therefore, we can write

$$\frac{Re_\theta}{M_e} = c\sqrt{x}\,, \tag{11}$$

for some constant c. Linearizing about the nominal value $x_{t,\mathrm{r}0}$ gives

$$\begin{aligned}\frac{Re_\theta}{M_e} &= c\sqrt{x_{t,\mathrm{r}0}} + \frac{c}{2\sqrt{x_{t,\mathrm{r}0}}}(x - x_{t,\mathrm{r}0})\,,\\ &= c\sqrt{x_{t,\mathrm{r}0}}\left(1 \pm \frac{x - x_{t,\mathrm{r}0}}{2x_{t,\mathrm{r}0}}\right)\,.\end{aligned} \tag{12}$$

This is in the same form as (10), thus the uncertainty φ is

$$\varphi = \pm\frac{x - x_{t,\mathrm{r}0}}{2x_{t,\mathrm{r}0}}\,. \tag{13}$$

Rearranging we can represent the transition location uncertainty as a function of φ

$$x_{t,\mathrm{r}} = x_{t,\mathrm{r}0}(1 \pm 2\varphi)\,, \tag{14}$$

thus the uncertainty range is $x_{t,\mathrm{r}} = x_{t,\mathrm{r}0} \cdot [0.6, 1.4]$ where $x_{t,\mathrm{r}0} = 145\mathrm{mm}$.

The same procedure is used to determine the uncertainty in $x_{t,\mathrm{c}}$. The nominal transition location is specified as 50mm downstream of the cowl leading edge, so the range is defined as $x_{t,\mathrm{c}} = x_{t,\mathrm{c}0} \cdot [0.6, 1.4]$ where $x_{t,\mathrm{c}0} = 50\mathrm{mm}$. The uncertainty ranges and nominal values for these parameters are summarized in Table 2.

Parameter	Min	Nominal	Max	Units
Stagnation Pressure	16.448	17.730	19.012	MPa
Stagnation Enthalpy	3.0551	3.2415	3.4280	MJ/kg
Angle of Attack	2.6	3.6	4.6	deg.
Turbulence Intensity	0.001	0.01	0.019	·
Turbulence Length Scale	0.1325	0.245	0.3575	m
Ramp Transition Location	0.087	0.145	0.203	m
Cowl Transition Location	0.030	0.050	0.070	m

Table 2: Summary of parameters and ranges used in the uncertainty quantification of the HyShot II scramjet within the 2D intake ramp simulation.

3.4. Quantity of interest

Numerical simulations where the shock system has moved all the way to the upstream boundary (i.e., when the flow is about to unstart) are computationally challenging. For this reason, we seek a proxy indicator for the unstart process. Several possible unstart indicators have been proposed [16]. Presently,

we use a function of the combustion chamber exit pressure. The motivation for this choice is related to the physics of compressible flow with heat addition (Rayleigh's flow): a supersonic flow decelerates towards sonic conditions when heat is released through a combustion process. The corresponding increase in pressure achieves its maximum just before choking conditions are reached.

The unstart proxy is the normalized integral of pressure over the last 10mm of the combustor

$$\frac{1}{\mathrm{vol}\,(\boldsymbol{V})} \iiint_{\boldsymbol{V}} P \; dx \; dy \; dz \,, \tag{15}$$

where $\boldsymbol{V}$ is the combustor region in $0.64 \leq x \leq 0.65$m. The integrated exit pressure (15) is a critical measure of the scramjet performance, and we treat it as the quantity of interest for uncertainty quantification.

4. Uncertainty quantification methodology

In this section, we use the notation $f(\mathbf{x})$ to represent the quantity of interest (integrated exit pressure defined in (15)) as a function of the input parameters (Table 2), where we treat the parameters as normalized to the hypercube $[-1, 1]^m$ ($m = 7$ in the HyShot II model). In other words, for a point $\mathbf{x} \in [-1, 1]^m$, the function $f(\mathbf{x})$ first shifts and linearly scales $\mathbf{x}$ to the application-specific ranges and then computes the quantity of interest from the given simulation with the proper inputs. This notation simplifies the methods' description.

The uncertainty quantification problem for the HyShot II model involves three distinct computations: (i) estimating lower and upper bounds for $f(\mathbf{x})$, (ii) identifying sets of inputs that produce safe operation defined as $f(\mathbf{x})$ remaining below a given threshold, and (iii) estimating a cumulative distribution function for $f(\mathbf{x})$ given a probability density function on $\mathbf{x}$. The last computation is the typical *forward* uncertainty quantification problem [54]. These computations are challenging for several reasons.

1. The exit pressure f is a functional of the pressure field computed from the complex, nonlinear, multiphysics simulation, so we have no prior knowledge of exploitable structure like linearity or convexity.
2. Each evaluation of $f(\mathbf{x})$ requires an expensive simulation; each run takes approximately 9500 CPU-hours on available computing resources.
3. Due to the multiphysics and domain coupling, we cannot evaluate gradients or Hessians of $f(\mathbf{x})$ with respect to $\mathbf{x}$.
4. Evaluations of $f(\mathbf{x})$ contain nonnegligible numerical error due to (i) the fixed point iteration scheme that solves the compressible flow model and (ii) the fixed mesh constructed to capture shocks at the nominal inputs.

Under these conditions, exhaustively sampling the seven-dimensional input space is infeasible for our three desired computations. Reducing the simulation's cost with parametric reduced-order models [1] is infeasible for the nonlinear, multiphysics simulation code. Multivariate approximation schemes for response

surfaces—e.g., splines or radial basis functions [61]—are precarious since we cannot densely sample the input space. Adaptive methods for response surface construction use heuristics (e.g., *expected improvement*) to increase efficiency [28], but their efficacy relies on a sufficiently accurate initial approximation, which is not guaranteed in high dimensions with limited samples. Sophisticated methods for response surfaces in high dimensions—e.g., sparse grids [4], low-rank approximations [13], or sparse approximation schemes [12]—are appropriate when $f(\mathbf{x})$ admits particular structure; we have no prior indication that the integrated exit pressure as a function of the HyShot II simulation inputs admits such structure. Global sensitivity metrics (e.g., Sobol indices) can identify the least important parameters, which we might fix at nominal values to reduce the dimension for the computations [50]. But such metrics require approximating seven-dimensional integrals; insufficient sampling can yield incorrect results.

4.1. Active subspaces and summary plots

Our approach discovers and exploits an *active subspace* in $f(\mathbf{x})$ [8]. The active subspace is defined by a set of important directions in the input space, which generalizes the idea of identifying the most important coordinates with sensitivity analysis. The important directions come from the first few eigenvectors of the $m \times m$ symmetric, positive semidefinite matrix

$$\boldsymbol{C} \;=\; \int \nabla f(\mathbf{x})\, \nabla f(\mathbf{x})\, \rho(\mathbf{x})\, d\mathbf{x} \;=\; \boldsymbol{W} \Lambda \boldsymbol{W}^T, \tag{16}$$

where ∇f is the gradient of f with respect to $\mathbf{x}$, ρ is a probability density function on the input space, $\boldsymbol{W}$ is the orthogonal matrix of eigenvectors, and Λ is the diagonal matrix of nonnegative eigenvalues. The ith eigenvalue λ_i satisfies

$$\lambda_i \;=\; \int (\mathbf{w}_i^T \nabla f(\mathbf{x}))^2\, \rho(\mathbf{x})\, d\mathbf{x}, \quad i = 1, \ldots, m. \tag{17}$$

In words, the eigenvalues measure how f changes, on average, in response to small perturbations in $\mathbf{x}$ along the corresponding eigenvectors. For $n < m$, if there is a large gap between the nth and $n+1$th eigenvalues, then a reasonable approximation for $f(\mathbf{x})$ is

$$f(\mathbf{x}) \;\approx\; g(\boldsymbol{W}_1^T \mathbf{x}), \tag{18}$$

where $\boldsymbol{W}_1$ contains the first n columns of $\boldsymbol{W}$, and g is a map from $\mathbb{R}^n$ to $\mathbb{R}$. The model defined by (18) is a low-dimensional approximation that can be exploited to enable otherwise infeasible parameter studies [8]. If the gradient $\nabla f(\mathbf{x})$ is available as a subroutine, then one can estimate $\boldsymbol{C}$ and its eigenpairs in (16) with Monte Carlo [9].

Our simulation code does not have a subroutine to compute the gradient. We do not trust finite difference approximations because of fixed point solver oscillations at nearby parameter values. To find the active subspace, we turn to methods for *sufficient dimension reduction* (SDR) developed in the context

of regression modeling [11]. The simplest of these methods is based on a least-squares-fit linear approximation of $f(\mathbf{x})$. Assume that $\rho : \mathbb{R}^m \to \mathbb{R}_+$ is a given probability density function on the inputs $\mathbf{x}$; as discussed in Section 3, we follow the maximum entropy principle to choose a uniform density on the hypercube $[-1, 1]^m$ in the absence of information to the contrary. Consider the procedure outlined in Algorithm 1.

Algorithm 1.
Active subspace estimation with least-squares linear approximation

1. Draw M samples $\{\mathbf{x}_j\}$ independently according to the density $\rho(\mathbf{x})$.
2. For each $\mathbf{x}_j$, compute $f_j = f(\mathbf{x}_j)$.
3. Determine the parameters $\hat{u}_0, \ldots, \hat{u}_m$ of the linear approximation

$$f(\mathbf{x}) \;\approx\; \hat{u}_0 + \hat{u}_1 x_1 + \cdots + \hat{u}_m x_m \tag{19}$$

   with least-squares

$$\hat{\mathbf{u}} \;=\; \underset{\mathbf{u}}{\operatorname{argmin}} \; \frac{1}{2} \|\boldsymbol{A}\mathbf{u} - \mathbf{f}\|_2^2, \tag{20}$$

   where

$$\boldsymbol{A} = \begin{bmatrix} 1 & \mathbf{x}_1^T \\ \vdots & \vdots \\ 1 & \mathbf{x}_M^T \end{bmatrix}, \quad \mathbf{f} = \begin{bmatrix} f_1 \\ \vdots \\ f_M \end{bmatrix}, \quad \hat{\mathbf{u}} = \begin{bmatrix} \hat{u}_0 \\ \vdots \\ \hat{u}_m \end{bmatrix}. \tag{21}$$

4. Compute the vector $\mathbf{w}$ as

$$\mathbf{w} \;=\; \frac{\hat{\mathbf{u}}'}{\|\hat{\mathbf{u}}'\|}, \tag{22}$$

   where $\hat{\mathbf{u}}' = [\hat{u}_1, \ldots, \hat{u}_m]^T$ is the last m coefficients (i.e., the gradient) of the linear approximation (19).

We define the active subspace with the unit vector $\mathbf{w}$ from (22), which identifies one important direction in the input space. The number M of samples must be $\mathcal{O}(m)$— large enough to fit the linear approximation (19) with least-squares—and the sample points $\{\mathbf{x}_j\}$ must produce a full rank least-squares problem. Some recent work studies the scaling of M with the dimension of the polynomial space for optimal approximation of the coefficients [41]; for a linear approximation (19), the polynomial space dimension is $m + 1$.

Several remarks are in order. We do not use the linear approximation as a predictive response surface. The coefficients $\hat{u}_1, \ldots, \hat{u}_m$ of the linear approximation identify a direction in the input space—hence the normalization in (22). The vector $\mathbf{w}$ may reveal low-dimensional structure in $f(\mathbf{x})$; our goal is to exploit this structure, if present, to enable otherwise infeasible uncertainty quantification with the expensive computer simulation.

In the SDR literature, Li and Duan [37, Theorem 2.1] show that under certain conditions, the least-squares-based procedure in Algorithm 1 identifies the dimension reduction space for a regression model. However, using the uniform density on the hypercube violates the assumptions of their theory. Additionally, treating the deterministic $f(\mathbf{x})$ with regression tools poses a challenge for interpretation; the standard regression model assumes random noise in the f_j's that is not present in our deterministic simulation. Nevertheless, we can apply the procedure as outlined and validate that $\mathbf{w}$ identifies the low-dimensional structure with a *summary plot* [11].

The summary plot is a scatter plot of $\mathbf{w}^T\mathbf{x}_j$ versus f_j. The plot may reveal a nearly one-to-one mapping between the linear combination $\mathbf{w}^T\mathbf{x}$ and the output $f(\mathbf{x})$. In other words, the output may be well-represented by a scalar-valued function of the linear combination of $\mathbf{x}$. Any perceived departure from the one-to-one mapping comes from one of two sources: (i) $f(\mathbf{x})$ may vary significantly as $\mathbf{x}$ moves orthogonally to $\mathbf{w}$ or (ii) the number M of samples may be too small to accurately estimate $\mathbf{w}$. In practice, one can determine whether the perceived departure is due to (i) or (ii) by running more simulations (i.e., increasing M).

If the f_j's depart very little from a one-to-one mapping, then we can reasonably approximate

$$f(\mathbf{x}) \;\approx\; g(\mathbf{w}^T\mathbf{x}), \quad \mathbf{x} \in [-1,1]^m, \tag{23}$$

where $g : \mathbb{R} \to \mathbb{R}$ is a function constructed from the pairs $(\mathbf{w}^T\mathbf{x}_j, f_j)$. This construction is similar to a single-index regression model for the exit pressure with a quadratic polynomial link function [25, Chapter 11]. Several constructions for g are possible—e.g., polynomial approximation or radial basis approximation—but we prefer not to exactly interpolate the f_j's due to finite sampling noise in the computed $\mathbf{w}$ from (22). Note the similarity between (23) and (18); both approximate $f(\mathbf{x})$ with a function of less than m linear combinations of $\mathbf{x}$.

The summary plot is a subjective tool, in general. The scientist perceives the f_j's departure from a supposed univariate function of $\mathbf{w}^T\mathbf{x}$, and she must decide if this departure is small enough to justify the model (23). One might use metrics such as the residual norm to compare different g's (e.g., different degrees of polynomial approximation), but such metrics can be deceiving without the aid of visualization. We use a bootstrap method to gain some confidence in the perceived structure.

4.2. Bootstrap for $\mathbf{w}$

The elements of $\mathbf{w}$ depend on the samples f_j used to fit the linear approximation (19). With M random samples, it is natural to ask if the gradient of the linear approximation has been sufficiently resolved to reveal the true one-dimensional active subspace, if one is present. But a limited budget of function evaluations prohibits us from checking convergence of $\mathbf{w}$ as M increases. We use a bootstrap technique [14] to estimate the variability in the computed components of $\mathbf{w}$ from the set of M samples.

Algorithm 2.
Bootstrap for estimating variability in $\mathbf{w}$

1. Choose the number N of *bootstrap replicates*.
2. For k from 1 to N, do the following.
   (a) Let
   $$\boldsymbol{\pi}^k \;=\; \begin{bmatrix}\pi_1^k & \cdots & \pi_M^k\end{bmatrix} \tag{24}$$
   be M-vectors of integers between 1 and M sampled with replacement.
   (b) Let the M-vector $\mathbf{f}_k$ and $M \times (m+1)$ matrix $\boldsymbol{A}_k$ be
   $$\mathbf{f}_k \;=\; \begin{bmatrix} f_{\pi_1^k} \\ \vdots \\ f_{\pi_M^k} \end{bmatrix}, \qquad \boldsymbol{A}_k \;=\; \begin{bmatrix} 1 & \mathbf{x}_{\pi_1^k}^T \\ \vdots & \vdots \\ 1 & \mathbf{x}_{\pi_M^k}^T \end{bmatrix}. \tag{25}$$
   (c) Compute $\mathbf{w}_k$ for each $\boldsymbol{A}_k$ and $\mathbf{f}_k$ as in (20) and (22).

The procedure outlined in Algorithm 2 samples, with replacement, the rows of $\boldsymbol{A}$ and elements of $\mathbf{f}$ in (20) to create a collection of bootstrap replicates $\{\mathbf{w}_k\}$. We can use these vectors to study the variability in the estimate $\mathbf{w}$ from (22)—e.g., with histograms that estimate the marginal bootstrap densities of $\mathbf{w}$'s components. Sharp peaks and narrow supports in the histograms provide confidence in $\mathbf{w}$. Wide supports may indicate a poor estimate of $\mathbf{w}$, or it may reflect insufficient sampling of the f_j's used to fit the linear approximation.

5. Discovering and exploiting a one-dimensional active subspace in the HyShot II quantity of interest

We apply Algorithm 1 to the exit pressure (15) as a function of the $m = 7$ input parameters in Table 2. We repeat the study for two values of the fuel plenum pressure $P_{0,\mathrm{H2}}$, specifically $P_{0,\mathrm{H2}} = 4.8$ bar (corresponding to a fuel/air equivalence ratio of 0.30 at nominal inflow conditions) and $P_{0,\mathrm{H2}} = 5.6$ bar (corresponding to a fuel/air equivalence ratio of 0.35 at nominal inflow conditions). The first is the nominal operating condition. The second is chosen to be close to the critical equivalence ratio of 0.38-to-0.39. Recent studies found a change in the flow regime from an *as-designed* fully supersonic combustion ($\phi <\sim 0.39$) to a state with a shock-train in the downstream part of the combustor (for $\phi >\sim 0.39$) [32, 34]. From an engineering point-of-view, it is interesting to see whether the effects of the inflow uncertainties change as one approaches this regime boundary. For $P_{0,\mathrm{H2}} = 4.8$ bar, we use $M = 50$ samples to estimate $\mathbf{w}$, which is roughly m^2. For $P_{0,\mathrm{H2}} = 5.6$ bar, we use $M = 14$ samples, which is $2m$. In both cases, the cost of searching for the one-dimensional active subspace is much less than the cost of estimating the integral quantities needed for global sensitivity analysis.

Figure 7 shows the summary plot for our application. The horizontal axis is the *active variable* $\mathbf{w}^T\mathbf{x}$, and the vertical axis is the exit pressure. The black circles are the samples f_j plotted against the corresponding value of the active variable $\mathbf{w}^T\mathbf{x}_j$. In our judgment, the circles indicate that we can approximate

the exit pressure with sufficient accuracy using a univariate function of $\mathbf{w}^T\mathbf{x}$ as in (23) for both $P_{0,\mathrm{H2}} = 4.8$ bar and $P_{0,\mathrm{H2}} = 5.6$ bar. We exploit this low-dimensional structure for the three desired computations needed for uncertainty quantification.

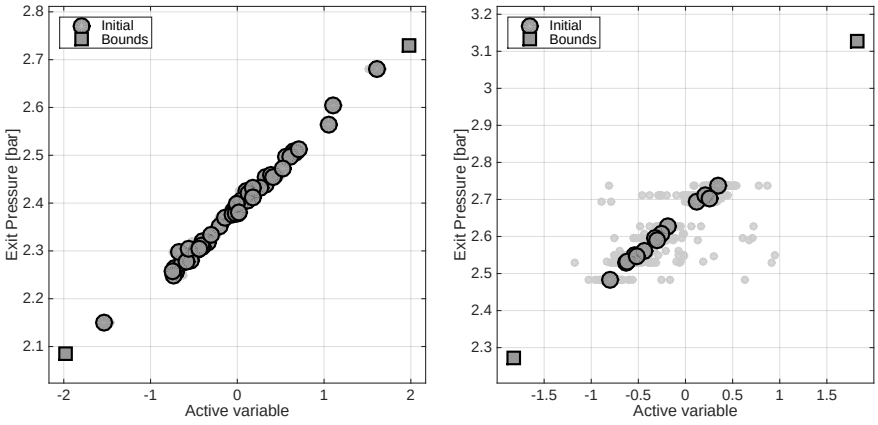

(a) $P_{0,\mathrm{H2}} = 4.8$ bar (b) $P_{0,\mathrm{H2}} = 5.6$ bar

Figure 7: The black circles are exit pressures computed from a set of HyShot II simulations and plotted against the corresponding active variable values $\mathbf{w}^T\mathbf{x}_j$. The left subfigure shows $M = 50$ runs with fuel plenum pressure $P_{0,\mathrm{H2}} = 4.8$ bar. The right subfigure shows $M = 14$ runs with $P_{0,\mathrm{H2}} = 5.6$ bar. A clear univariate relationship exists for both cases that we exploit to quantify uncertainty. The overlapping gray circles distributed horizontally correspond to exit pressures plotted against 100 bootstrap replicates of the active variable. There is hardly any variability in the bootstrap estimates for $P_{0,\mathrm{H2}} = 4.8$. The horizontal variability for $P_{0,\mathrm{H2}} = 5.6$ bar is a result of the relatively low oversampling ($M = 2m$) when fitting the linear approximation (19). The gray squares show the exit pressures at the boundaries of the domain where the perceived trend suggests we find the upper and lower bounds of the exit pressure. The univariate relationship is validated by these two additional simulations for each $P_{0,\mathrm{H2}}$.

5.1. Sensitivity analysis

The components of $\mathbf{w}$ measure the global sensitivity of the exit pressure to each of the seven parameters; see [50, Section 1.2.5]. These values and their corresponding input parameters are shown in Table 3. They suggest that four of the seven parameters contribute the most to the one-dimensional active subspace. The angle of attack, the stagnation conditions (temperature and pressure), and the turbulence intensity dominate the transition location and the turbulence length scale. The largest contributors to changes in the exit pressure are the inputs related to the intensity of the bow shock at the nose of the vehicle, which influences the thermodynamic post shock conditions and the flow entering the Hyshot combustor. The state of the boundary layers (as determined by different transition location) plays a secondary role. It is more difficult to assess why the turbulence intensity plays such an important role. Caution is needed before

interpreting this result, since RANS eddy-viscosity models (like the SST) are known to behave erratically across strong shock waves—like those present here [52]; a more detailed investigation is outside the scope of the present paper.

There is an important distinction between the plenum pressures $P_{0,\mathrm{H2}} = 4.8$ bar and 5.6 bar. The leading factors in determining changes in the exit pressure do not change, but the relative importance of stagnation pressure and enthalpy does. Large-eddy simulations of the HyShot II combustor [32] have shown that the combustion is close to complete, and thus the total heat release is essentially set by the fuel mass flux and thus independent of variations in the inflow conditions. Those same large-eddy simulations also showed that the heat losses through the walls and the momentum loss due to friction are very large in this combustor. Therefore, the most plausible explanation for the sensitivity of the exit pressure to the inflow stagnation pressure and enthalpy is that these variations must alter the friction and/or wall heat losses.

Index	$P_{0,\mathrm{H2}} = 4.8$ bar	$P_{0,\mathrm{H2}} = 5.6$ bar	Parameter
1	0.6506	0.7066	Angle of Attack
2	0.5565	0.5008	Turbulence Intensity
3	-0.0002	0.0289	Turbulence Length Scale
4	0.3685	0.2051	Stagnation Pressure
5	-0.3566	-0.4490	Stagnation Enthalpy
6	-0.0196	-0.0591	Cowl Transition Location
7	0.0607	0.0432	Ramp Transition Location

Table 3: The components of the vector $\mathbf{w}$ that defines the one-dimensional active subspace—computed independently for each $P_{0,\mathrm{H2}}$. Each component corresponds to one of the parameters in the HyShot II simulation as described in Section 3.

5.2. Bootstrap results

We apply the bootstrap procedure from Algorithm 2 to $\mathbf{w}$ for the HyShot II exit pressure using $N = 100$ bootstrap replicates. Note that the bootstrap uses only the available simulation runs, so its cost is negligible. The bootstrap histograms along with stem plots of the components of $\mathbf{w}$ from Table 3 are shown for $P_{0,\mathrm{H2}} = 4.8$ bar in Figure 8 and for $P_{0,\mathrm{H2}} = 5.6$ bar in Figure 9. The histograms' sharp peaks around the stems suggests confidence in the computed directions. The relatively large ranges in Figure 9 are a result of the low oversampling factor ($M = 2m$) used to fit the linear approximation; with too few samples to draw from, the probability is greater that a bootstrap replicate is not representative of the true values of $\mathbf{w}$.

We can study how bootstrap variability in $\mathbf{w}$ affects the perceived relationship in the summary plot. We use the 100 bootstrap replicates $\mathbf{w}_k$ to plot the exit pressures f_j against the corresponding active variables $\mathbf{w}_k^T\mathbf{x}_j$. The result is a horizontal scatter of gray dots around each point in Figure 7. The scatter provides a visual indication of how $\mathbf{w}$'s variability affects the perceived relationship between the active variable $\mathbf{w}^T\mathbf{x}$ and the exit pressure f. There is hardly any

scatter in the case of $P_{0,\mathrm{H2}} = 4.8$ bar with $M = 50$ samples; the 100 gray circles per original sample (black circle) are barely visible behind the original samples. For $P_{0,\mathrm{H2}} = 5.6$ bar, the large spread relative to the range of the active variable is due to the low sampling ($M = 2m$) used to fit the linear approximation.

In the lower fuel pressure case, the bootstrap instills confidence as intended, because 50 samples is sufficient to produce meaningful results with resampling. In the high fuel pressure case, 14 samples is not enough for the bootstrap to instill confidence—although 14 samples in the summary plot (Figure 7(b)) still show a strong univariate relationship between the active variable and the exit pressure. Despite the bootstrap variability in the second case, three things provide confidence of the one-dimensional active subspace: (i) the strong validation in the first case with 50 runs and low bootstrap variability, (ii) the physical intuition that the relationship between inputs and outputs should not change dramatically between the two cases, and (iii) the apparent relationship in the summary plot Figure 7(b).

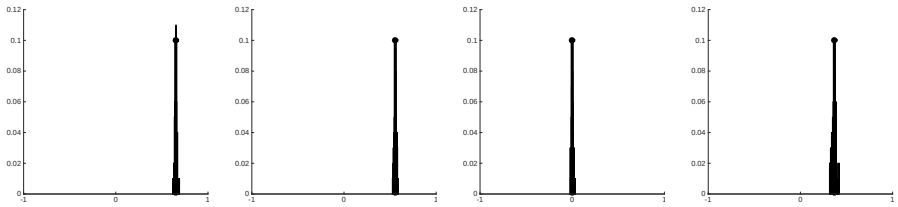

(a) Angle of Attack (b) Turb. Intensity (c) Turb. Len. Scale (d) Stag. Pres.

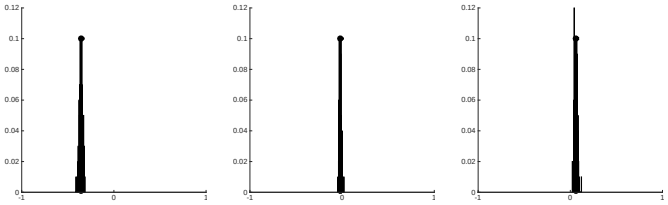

(e) Stag. Enth. (f) Cowl Trans. (g) Ramp Trans.

Figure 8: Bootstrap histograms of the components of the active subspace vector $\mathbf{w}$ for $P_{0,\mathrm{H2}} = 4.8$ bar. The black stems are the computed components of $\mathbf{w}$ from Table 3. The caption of each subfigure names the specific inflow parameter. The sharp peaks around each of the stems provides confidence that the computed $\mathbf{w}$ are stable.

We exploit the one-dimensional active subspace revealed in the summary plots (Figure 7) and confirmed by the bootstrap to quantify uncertainty in the HyShot II's simulated prediction of exit pressure. The uncertainty quantification is comprised of three computations: (i) estimating the range of possible exit pressures, (ii) identifying safe operating conditions, and (iii) estimating a cumulative distribution function for the exit pressure.

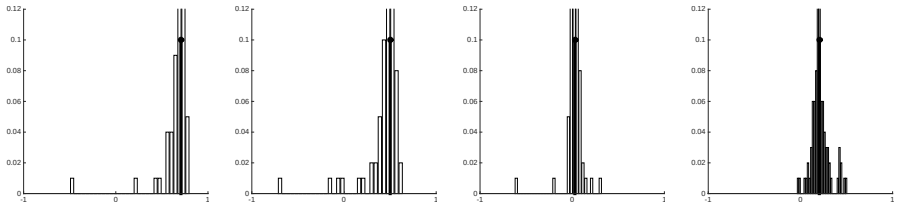

(a) Angle of Attack (b) Turb. Intensity (c) Turb. Len. Scale (d) Stag. Pres.

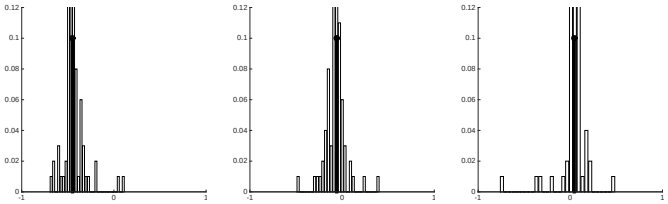

(e) Stag. Enth. (f) Cowl Trans. (g) Ramp Trans.

Figure 9: Bootstrap histograms of the components of the active subspace vector $\mathbf{w}$ for $P_{0,\mathrm{H2}} = 5.6$ bar. The black stems are the computed components of $\mathbf{w}$ from Table 3. The caption of each subfigure names the specific inflow parameter. The sharp peaks around each of the stems provides confidence that the computed $\mathbf{w}$ are stable. The relatively large spread compared to Figure 8 is due to the relatively small number ($M = 14$) of samples.

5.3. Approximating the range of exit pressures

The conditions noted at the beginning of Section 4 are common in design optimization problems with expensive computer simulations. Several approaches attempt to optimize with the aid of a response surface [2, 28, 63]. Unfortunately, our simulation is too expensive to use the iterative procedures that refine in possible regions of the optima.

The summary plots in Figure 7 suggest that the exit pressure $f(\mathbf{x})$ is a monotonic function of the active variable $\mathbf{w}^T\mathbf{x}$. A univariate, monotonic, and continuous function defined on a bounded interval can be bounded by the function values at the interval boundaries. We use the following heuristic to estimate the maximum and minimum exit pressures. Define

$$\mathbf{x}_{\max} = \underset{-1\le\mathbf{x}\le 1}{\operatorname{argmax}}\ \mathbf{w}^T\mathbf{x}, \qquad f_{\max} = f(\mathbf{x}_{\max}). \tag{26}$$

Similarly define

$$\mathbf{x}_{\min} = \underset{-1\le\mathbf{x}\le 1}{\operatorname{argmin}}\ \mathbf{w}^T\mathbf{x}, \qquad f_{\min} = f(\mathbf{x}_{\min}). \tag{27}$$

The components of $\mathbf{x}_{\max}$ and $\mathbf{x}_{\min}$ are either -1 or 1, and the two points correspond to opposite corners of the hypercube, i.e., $\mathbf{x}_{\min} = -\mathbf{x}_{\max}$. The signs of $\mathbf{x}_{\max}$'s components are determined entirely by the corresponding signs of $\mathbf{w}$'s components.

The range $[f_{\min}, f_{\max}]$ provides our estimate for the range of exit pressures from the HyShot II model. Computing this interval requires running the model four times beyond the initial set of samples used to compute $\mathbf{w}$—two for each $P_{0,\mathrm{H2}}$—which is much cheaper than adaptively constructing a response surface. The values of $f_{\min}$ and $f_{\max}$ are shown in Figure 7 as black squares along with the samples (black circles) used to compute $\mathbf{w}$ defining the active subspace—all plotted against the active variable. These runs confirm the monotonic structure perceived in the summary plot. At worst, these runs bound the initial exit pressure samples. At best, they provide estimates of the range of exit pressures over all values of the input parameters. Admittedly, this approach is not far from *guess-and-check*. But the guesses were informed by the structure revealed in the summary plot. Checking the necessary conditions for stationarity of these points is not feasible since we do not have first or second derivatives. One can interpret this heuristic as one step of a least-squares-fit linear approximation-based approach for derivative-free optimization with one global linear model [7]. However, the important difference is that the summary plot provides evidence that the one global linear model is sufficient to estimate the optima.

It is natural to wonder if the estimated range $[f_{\min}, f_{\max}]$ is strongly affected by the randomness in $\mathbf{w}$ originating in the random sampling in Algorithm 1. If some of $\mathbf{w}$'s components are sufficiently far from zero, then small perturbations in $\mathbf{w}$ do not change the signs of the corresponding components of $\mathbf{x}_{\min}$ and $\mathbf{x}_{\max}$. If a component of $\mathbf{w}$ is small in magnitude—such as components 2, 6, and 7 in Table 3—then small, random changes may change the corresponding components of $\mathbf{x}_{\min}$ and $\mathbf{x}_{\max}$ from -1 to 1 or vice versa. However, by the active subspace's construction, changes in these parameters do not change the exit pressure as much as changes in parameters with large $\mathbf{w}$ components. This gives confidence that the estimated range is stable under small changes in $\mathbf{w}$.

5.4. Constraining the exit pressure

Consider the following exercise in safety engineering. Suppose that the scramjet operates safely when the exit pressure is below 2.8 bar, but it nears unsafe operation above 2.8 bar. With the one-dimensional active subspace, we can quickly characterize the parameter regime that produces exit pressures below the threshold of 2.8 bars.

The first step is to build a response surface model of the exit pressure as a function of the active variable $\mathbf{w}^T\mathbf{x}$. We could try to construct a response surface of all seven input parameters. But with only 52 model runs for $P_{0,\mathrm{H2}} = 4.8$ bar (50 runs to compute $\mathbf{w}$ and 2 runs to estimate the range of f) and 16 model runs for $P_{0,\mathrm{H2}} = 5.6$ bar (14 for $\mathbf{w}$ and 2 for the estimated range), our modeling choices would be very limited. The apparent low departure from a univariate relationship in Figure 7 suggests that we can construct a useful response of just the active variable as in (23). In particular, we can use the exit pressure samples f_j and corresponding active variable values $\mathbf{w}^T\mathbf{x}_j$ to model the exit pressure as a univariate quadratic polynomial of the active variable.

In statistical regression, the coefficient of determination (i.e., R^2) is one of several metrics for the quality of the model [60]. This coefficient is interpreted

as the proportion of variance in the response (the output) that comes from variance in the predictors (the inputs); R^2 values near 1 indicate a good model for the data. This interpretation is not appropriate in our case, since there is no random noise in the f_j's. Nevertheless, we report the R^2 values for the quadratic model of the active variable and treat them as a measure of discrepancy between the model and the data: $R^2 = 0.993$ for $P_{0,\mathrm{H2}} = 4.8$ bar and $R^2 = 0.998$ for $P_{0,\mathrm{H2}} = 5.6$ bar.

Regression predictions often include frequentist confidence bounds. We can compute the upper 99% confidence bound and treat it as a conservative factor when seeking the safe operating conditions.[1] We find the largest value of the active variable such that the upper confidence bound from the quadratic approximation is less than the threshold of 2.8 bars. Figure 10 shows the quadratic approximation and its upper 99% confidence limit for both values of $P_{0,\mathrm{H2}}$. The shaded region identifies the values of the active variable that produce exit pressures at or below the pressure threshold.

Let $y_{\max}$ be the value of the active variable where the upper confidence bound crosses the pressure threshold. Then the safe region of the normalized input parameters is the set $\mathcal{S}$ defined as

$$\mathcal{S} \;=\; \{\mathbf{x} \,:\, \mathbf{w}^T\mathbf{x} \leq y_{\max},\; -1 \leq \mathbf{x} \leq 1\}. \tag{28}$$

One can easily shift and scale this region to the space of the HyShot II model's input space for a physical interpretation. The linear inequality constraint implies that the inputs are related with respect to the exit pressure. For example, the range of safe angles of attack depends on the other input variables. The presence of the active subspace and the quality of the quadratic approximation enable us to simply characterize these relationships.

The safe set $\mathcal{S}$ defined in (28) is like a seven-dimensional box with the top chopped off by the hyperplane $\mathbf{x}^T\mathbf{w} \leq y_{\max}$ (though the notion of "top" is problematic in seven dimensions). We can identify a new set of independent ranges for the input variables such that all inputs within those ranges produce exit pressures below the 2.8 bar safety threshold—according to the quadratic response surface model. This is like finding the largest seven-dimensional box that fits inside the set $\mathcal{S}$. More precisely, we solve the following optimization problem,

$$\begin{array}{ll} \underset{\mathbf{x}}{\text{maximize}} & \prod_{i=1}^{m} |x_i - x_{i,\min}|, \\ \text{subject to} & \mathbf{x} \in \mathcal{S}, \end{array} \tag{29}$$

where $x_{i,\min}$ are the components of the minimizer $\mathbf{x}_{\min}$ from (27). The maximizing components define the corner of the largest hyperrectangle opposite the corner $\mathbf{x}_{\min}$. We can shift and scale this hyperrectangle to the space of model parameters for physical interpretation. Table 4 shows the ranges of parameters affected by the safety condition.

[1] Since there is no noise in the f_j's, the statistical interpretation of the upper confidence bound is not valid; for example, we cannot interpret the 99% confidence bounds as a random interval that contains the exit pressure with probability 0.99.

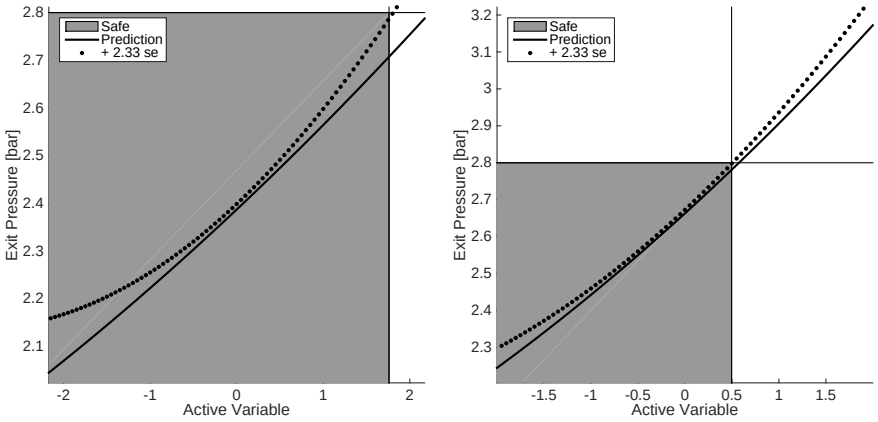

(a) $P_{0,\mathrm{H2}} = 4.8$ bar (b) $P_{0,\mathrm{H2}} = 5.6$ bar

Figure 10: A quadratic polynomial models the relationship between the active variable $\mathbf{w}^T\mathbf{x}$ and the exit pressure. The solid line shows the quadratic model. The dotted line shows the computed upper 99% confidence bound. We find the value of the active variable $y_{\max}$ where the upper confidence limit crosses the safety threshold of 2.8 bars. All values of the active variable less than $y_{\max}$ produce pressures within the safety limit. The set of safe input variables is shown in (28).

	Parameter	Min	Max	Units
$P_{0,\mathrm{H2}} = 4.8$ bar	Angle of Attack	2.6	4.29	deg.
	Turbulence Intensity	0.001	0.0188	·
$P_{0,\mathrm{H2}} = 5.6$ bar	Stagnation Enthalpy	3.15	3.4280	MJ/kg
	Angle of Attack	2.6	3.54	deg.
	Turbulence Intensity	0.001	0.013	·

Table 4: Updated ranges for input parameters that lead to safe scramjet operation with exit pressure, modeled with a quadratic function of the active variable, less than 2.8 bars. The parameters listed for each $P_{0,\mathrm{H2}}$ case are the parameters affected by the restriction in (28). Compare these ranges to those in Table 2.

We can interpret such analysis as backward uncertainty propagation that characterizes safe inputs given a characterization of a safe output under the constraint that the inputs be independent. These spaces differ between the two fuel plenum pressures; see Figure 10 and Table 4. For $P_{0,\mathrm{H2}} = 4.8$ bar, only angle of attack and turbulence intensity are affected by the safety constraint on the pressure; angle of attack must be less than 4.29 degrees, and turbulence intensity must be less than 0.0188. For $P_{0,\mathrm{H2}} = 5.6$ bar, the same parameters are constrained—angle of attack less than 3.54 and turbulent intensity less than 0.013—and a stricter minimum on stagnation enthalpy (greater than 3.15 MJ/Kg) appears. To explain this, we note that low free-stream stagnation temperature leads to conditions that are further away from the adiabatic frame temperature, which leads to increased heat deposition in the air-stream. This

in turns leads to more deceleration and larger pressure increases in the chamber [62].

5.5. Cumulative distribution function

We use the quadratic approximation of the active variable to estimate a cumulative distribution function of the exit pressure for both fuel plenum pressures. We draw 5000 samples independently and uniformly from the scramjet's seven-dimensional input space (i.e., according the the uniform density $\rho(\mathbf{x})$), and for each sample we evaluate (i) the active variable and (ii) the quadratic approximation of exit pressure as a function of the active variable. Figure 11 shows estimates of the cumulative distribution function computed from these samples with a Gaussian kernel density estimator. The vertical lines show the estimated upper and lower bounds computed in Section 5.3.

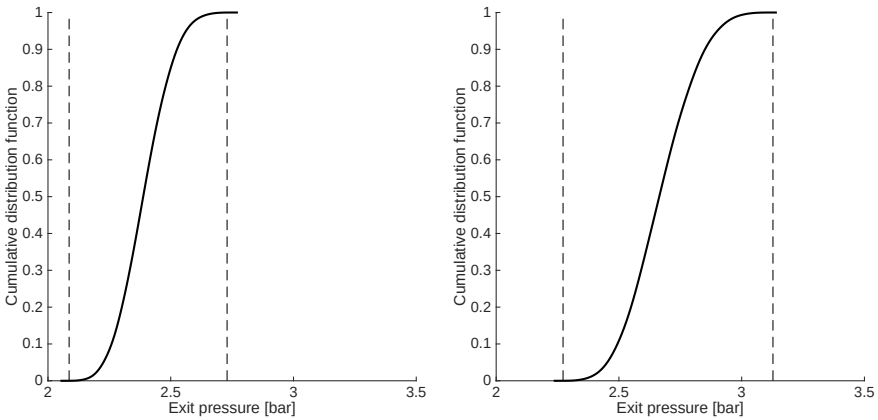

(a) $P_{0,\mathrm{H2}} = 4.8$ bar (b) $P_{0,\mathrm{H2}} = 5.6$ bar

Figure 11: Estimated cumulative distribution functions for exit pressure at both values of fuel plenum pressure $P_{0,\mathrm{H2}}$. These are estimated with Gaussian kernel density estimates; the samples are drawn from the quadratic approximation of exit pressure as a function of the active variable. Vertical lines show the estimated bounds for each case estimated in Section 5.3.

6. Summary and discussion

We present a numerical investigation of the reactive flow within a hydrogen-fueled scramjet with the objective of studying the effect of uncertainties in operating conditions on the overall performance. We carry out three-dimensional RANS simulations with a flamelet-based combustion model at two different fuel plenum pressures and record the pressure at the engine exit as a measure of system performance. We consider seven uncertain parameters; their ranges are justified by experimental evidence and/or expert opinions.

The uncertainty quantification starts by identifying a one-dimensional reparameterization of the map between simulation parameters and the exit pressure. This reparameterization is based on a one-dimensional active subspace; the coefficients of a global, least-squares-fit linear approximation define the active subspace and are validated by both a summary plot and a bootstrap. Fitting an accurate linear approximation for a function of m variables takes $\mathcal{O}(m)$ simulations. For the smaller fuel plenum pressure, we use 50 simulations to estimate the vector defining the active subspace; for the larger fuel plenum pressure, we use only 14 simulations. In both cases, the summary plot shows strong evidence of near one-to-one map between the linear combination of the inputs (i.e., the active variable) and the exit pressure. However, with only 14 runs, the bootstrap results show much greater variability due to insufficient sampling. We treat the components of the vector defining the active subspace as sensitivity metrics to gain insight into the physics.

With a one-dimensional approximation and the perceived monotonic relationship in the summary plot, the computations necessary to quantify uncertainty become much easier. The minimum and maximum values for a univariate, continuous, monotonic function reside at the endpoints of the interval domain; we exploit this fact to estimate the range of possible exit pressures over all parameter values for both fuel pressure cases. We can validate a univariate response surface with the summary plot to complement standard quality metrics (e.g., residual norms). We use the validated, univariate, monotonic response surface to identify safe operating conditions in the active variable, and these constraints translate to the original model parameters. Lastly, we use the validated response surface as a surrogate to estimate a cumulative distribution function of the exit pressure given variability in the input parameters. To quantify uncertainty for both fuel pressure cases, we used 68 full simulations: 64 to identify the respective one-dimensional active subspaces and 4 to estimate the respective exit pressure ranges.

Quantifying uncertainty in such a complex and expensive simulation with seven independent inputs using so few runs is remarkable. It was only possible because of the low-dimensional structure revealed by the active subspace. If the summary plots had not shown a univariate, monotonic relationship, then our simple heuristics for estimating the exit pressure range and validating a response surface would not be justified. Thus, our approaches are only appropriate in the presence of this type of structure. The innovation in active subspaces is the potential to cheaply identify such exploitable structure to make possible otherwise infeasible computations.

Acknowledgments

This research was funded by the U.S. Department of Energy [National Nuclear Security Administration] under Award No. NA28614. Additional computational resources were made possible by the following award, MRI-R2: Acquisition of a Hybrid CPU/GPU and Visualization Cluster for Multidisciplinary Studies in Transport Physics with Uncertainty Quantification. This award is

funded under the American Recovery and Reinvestment Act of 2009 (Public Law 111-5). The first author's work is supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC-0011077.

References

References

[1] Benner, P., Gugercin, S., Willcox, K., 2013. A survey of model reduction methods for parametric systems.

[2] Billups, S., Larson, J., Graf, P., 2013. Derivative-free optimization of expensive functions with computational error using weighted regression. SIAM Journal on Optimization 23 (1), 27–53.
URL http://dx.doi.org/10.1137/100814688

[3] Boyce, R., Gerard, S., Paull, A., 2014/08/14 2003. The HyShot Scramjet Flight Experiment - Flight Data and CFD Calculations Compared. American Institute of Aeronautics and Astronautics.
URL http://dx.doi.org/10.2514/6.2003-7029

[4] Bungartz, H.-J., Griebel, M., 5 2004. Sparse grids. Acta Numerica 13, 147–269.
URL http://journals.cambridge.org/article_S0962492904000182

[5] Chapuis, M., Fedina, E., Fureby, C., Hannemann, K., Karl, S., Schramm, J., 2013. A computational study of the HyShot II combustor performance. In: Proceedings of the Combustion Institute. Vol. 34. pp. 2101–2109.

[6] Chen, C., Liu, D., 2011. Numerical investigation of supersonic combustion of the HyShot II in the shock tunnel. Journal of Aeronautics, Astronautics and Aviation 43 (2), 119–128.

[7] Conn, A. R., Scheinberg, K., Vicente, L. N., 2009. Introduction to Derivative-free Optimization. SIAM, Philadelphia.

[8] Constantine, P., Dow, E., Wang, Q., 2014. Active subspace methods in theory and practice: applications to kriging surfaces. SIAM Journal on Scientific Computing 36 (4), A1500–A1524.
URL http://dx.doi.org/10.1137/130916138

[9] Constantine, P., Gleich, D., 2015. Computing active subspaces with Monte Carlo. arXiv preprint arXiv:1408.0545v2.

[10] Constantine, P. G., Emory, M., Palacios, F., Kseib, N., Iaccarino, G., 2013. Quantification of margins and uncertainties using an active subspace method for approximating bounds. In: 11th International Conference on Structural Safety & Reliability.

[11] Cook, R. D., 2009. Regression Graphics: Ideas for Studying Regressions through Graphics. John Wiley & Sons.

[12] Doostan, A., Owhadi, H., 2011. A non-adapted sparse approximation of PDEs with stochastic inputs. Journal of Computational Physics 230 (8), 3015 – 3034.
URL http://www.sciencedirect.com/science/article/pii/S0021999111000106

[13] Doostan, A., Validi, A., Iaccarino, G., 2013. Non-intrusive low-rank separated approximation of high-dimensional stochastic models. Computer Methods in Applied Mechanics and Engineering 263 (0), 42 – 55.
URL http://www.sciencedirect.com/science/article/pii/S0045782513000935

[14] Efron, B., Tibshirani, R., 1993. An Introduction to the Bootstrap. Chapman & Hall/CRC.

[15] Emory, M., Larsson, J., Iaccarino, G., 2013. Modeling of structural uncertainties in Reynolds-averaged Navier-Stokes closures. Physics of Fluids 25 (110822).
URL http://dx.doi.org/10.1063/1.4824659

[16] Emory, M., Terrapon, V., Pečnik, R., Iaccarino, G., April 2011. Characterizing the operability limits of the HyShot II scramjet through RANS simulations. In: 17^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2011-2282.
URL http://arc.aiaa.org/doi/pdf/10.2514/6.2011-2282

[17] Fureby, C., Chapuis, M., Fedina, E., Karl, S., 2011. CFD analysis of the HyShot II scramjet combustor. Proc. Comb. Inst. 33, 2399–2405.

[18] Gardner, A., December 2004. HyShot scramjet testing in the HEG. Ph.D. thesis, University of Queensland.

[19] Gardner, A., Hannemann, K., Steelant, J., Paull, A., July 2004. Ground testing of the HyShot supersonic combustion flight experiment in HEG and comparison with flight data. In: 40^{th} AIAA Joint Propulsion Conference. No. AIAA-2004-3345.

[20] German Aerospace Center, D., 2012. DLR Report to Stanford. Document provided to Stanford by DLR for HyShot II simulation collaboration.

[21] Ghanem, R. G., Spanos, P. D., 1991. Stochastic Finite Elements: A Spectral Approach. Springer.

[22] Hannemann, K., Karl, S., Schramm, J., Steelant, J., 2010. Methodology of a combined ground based testing and numerical modeling analysis of supersonic combustion flow paths. Shock Waves 20 (5), 353–366.

[23] Hannemann, K., Martinez Schramm, J., Karl, S., Steelant, J., 2009. Experimental investigation of different scramjet hydrogen injection systems. In: Proc. 6th European Symposium on Aerothermodynamics for Space Vehicles. ESA-SP-659.

[24] Hass, N., Smart, M., Paull, A., 2005. Flight data analysis of HyShot II. In: 13^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2005-3354.

[25] Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning, 2nd Edition. Springer.

[26] Hong, Z., Davidson, D., Hanson, R., 2011. An improved H2/O2 mechanism based on recent shock tube/laser absorption measurements. Combustion and Flame 158, 633–644.

[27] Jaynes, E. T., 1957. Information theory and statistical mechanics. Physical review 106 (4), 620.

[28] Jones, D. R., 2001. A taxonomy of global optimization methods based on response surfaces. Journal of global optimization 21 (4), 345–383.

[29] Karl, S., Hannemann, K., Mack, A., Steelant, J., 2008. CFD analysis of the HyShot II scramjet experiments in the HEG shock tunnel. In: 15^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2008-2548.

[30] Karl, S., Schramm, J., Laurence, S., Hannemann, K., April 2011. CFD analysis of unstart characteristics of the HyShot II scramjet configuration in the HEG shock tunnel. In: 17^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2011-2309.

[31] Koehler, J., Owen, A., 1996. Computer experiments. In: Ghosh, S., Rao, C. (Eds.), Design and Analysis of Experiments. Vol. 13 of Handbook of Statistics. Elsevier, pp. 261 – 308.
URL http://www.sciencedirect.com/science/article/pii/S016971619613011X

[32] Larsson, J., Laurence, S. J., Bermejo-Moreno, I., Bodart, J., Karl, S., Vicquelin, R., 2014. Incipient thermal choking and stable shock-train formation in the heat-release region of a scramjet combustor. Part II: Large eddy simulations. Combust. Flame, submitted.

[33] Laurence, S. J., Karl, S., Hannemann, K., 2013. Experimental and numerical investigation of the HyShot II flight experiment. In: ISSW30. Madison, USA.

[34] Laurence, S. J., Lieber, D., Martinez Schramm, J., Hannemann, K., Larsson, J., 2014. Incipient thermal choking and stable shock-train formation in the heat-release region of a scramjet combustor. Part I: Shock-tunnel experiments. Combust. Flame, accepted.

[35] Laurence, S. J., Martinez Schramm, J., Karl, S., Hannemann, K., 2011. An experimental investigation of steady and unsteady combustion phenomena in the HyShot II combustor. AIAA Paper 2011-2310.

[36] Laurence, S. J., Ozawa, H., Lieber, D., Martinez Schramm, J., Hannemann, K., 2012. Investigation of unsteady/quasi-steady scramjet behavior using high-speed visualization techniques. AIAA Paper 2012-5913.

[37] Li, K.-C., Duan, N., 1989. Regression analysis under link violation. The Annals of Statistics 17 (3), pp. 1009–1052.
URL http://www.jstor.org/stable/2241708

[38] Menter, F., 1993. Zonal Two Equation k-w Turbulence Models For Aerodynamic Flows. American Institute of Aeronautics and Astronautics.
URL http://dx.doi.org/10.2514/6.1993-2906

[39] Menter, F. R., 1994. Two-equation eddy-viscosity turbulence models for engineering applications. AIAA journal 32 (8), 1598–1605.

[40] Menter, F. R., Langtry, R., Völker, S., 2004. Transition modeling for general purpose CFD codes. Flow, Turbulence, and Combustion 77, 277–303.

[41] Migliorati, G., Nobile, F., von Schwerin, E., Tempone, R., 2014. Analysis of discrete l^2 projection on polynomial spaces with random evaluations. Foundations of Computational Mathematics 14 (3), 419–456.
URL http://dx.doi.org/10.1007/s10208-013-9186-4

[42] O'Byrne, S., Doolan, M., Olsen, S. R., Houwing, A. F. P., 2000. Analysis of transient thermal choking processes in a model scramjet engine. Journal of Propulsion and Power 16 (5), 808–814.

[43] Oliver, A., Lillard, R., Schwing, A., Blaisdell, G., Lyrintzis, A., et al., 2007. Assessment of turbulent shock-boundary layer interaction computations using the overflow code. AIAA Paper 104.

[44] Owen, A. B., 2013. Monte Carlo theory, methods and examples.

[45] Pečnik, R., Terrapon, V., Ham, F., Iaccarino, G., 2010. Full system RANS of the HyShot II scramjet part 1: numerics and non-reactive cases. Annual research briefs, Center for Turbulence Research.

[46] Pečnik, R., Terrapon, V., Ham, F., Iaccarino, G., Pitsch, H., Aug 2012. Reynolds-averaged Navier-Stokes simulations of the HyShot II scramjet. AIAA Journal 50 (8), 1717–1732.

[47] Pečnik, R., Witteveen, J. A., Iaccarino, G., 2013. Assessment of uncertainties in modeling of laminar to turbulent transition for transonic flows. Flow, Turbulence and Combustion 91 (1), 41–61.
URL http://dx.doi.org/10.1007/s10494-013-9453-0

[48] Rasmussen, C. E., Williams, C. K., 2006. Gaussian Processes for Machine Learning. The MIT Press.

[49] Saghafian, A., Terrapon, V., Ham, F., Pitsch, H., 2011. An efficient flamelet-based combustion model for supersonic flows. In: 17^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2011-2267.

[50] Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. Global Sensitivity Analysis: The Primer. John Wiley & Sons.

[51] Schramm, J. M., Karl, S., Hannemann, K., Steelant, J., 2008. Ground testing of the HyShot II scramjet configuration in HEG. In: 15^{th} AIAA International Space Planes and Hypersonic Systems and Technologies Conference. No. AIAA-2008-2547.

[52] Sinha, K., Mahesh, K., Candler, G. V., 2003. Modeling shock unsteadiness in shock/turbulence interaction. Phys. Fluids 15 (8), 2290–2297.

[53] Smart, M., Hass, N., Paull, A., 2006. Flight data analysis of the HyShot II scramjet flight experiment. AIAA Journal 44 (10), 2366–2375.

[54] Smith, R. C., 2013. Uncertainty Quantification: Theory, Implementation, and Applications. SIAM, Philadelphia.

[55] Terrapon, V., Pečnik, R., Ham, F., Iaccarino, G., 2010. Full system RANS of the HyShot II scramjet part 2: reactive cases. Annual research briefs, Center for Turbulence Research.

[56] Terrapon, V., Pečnik, R., Ham, F., Pitsch, H., 2009. A flamelet-based model for supersonic combustion. Annual research brief, Center for Turbulence Research.

[57] Tu, J., Yeoh, G. H., Liu, C., 2008. Computational Fluid Dynamics: A Practical Approach. Elsevier, Inc.

[58] Wagner, J. L., Yuceil, K. B., Valdivia, A., Clemens, N. T., Dolling, D. S., 2009. Experimental investigation of unstart in an inlet/isolator model in Mach 5 flow. AIAA Journal 47 (6), 1528–1542.

[59] Waltrup, P. J., Billig, F. S., 1973. Prediction of precombustion wall pressure distributions in scramjet engines. Journal of Spacecraft and Rockets 10 (9), 620–622.

[60] Weisberg, S., 2005. Applied Linear Regression. John Wiley & Sons.

[61] Wendland, H., 2004. Scattered Data Approximation. Cambridge University Press.

[62] Wendt, M. N., Stalker, R. J., Jacobs, P. A., 2014/08/17 1997. Fuel stagnation temperature effects on mixing with supersonic combustion flows. Journal of Propulsion and Power 13 (2), 274–280.
URL http://dx.doi.org/10.2514/2.5159

[63] Wild, S. M., Regis, R. G., Shoemaker, C. A., 2008. ORBIT: Optimization by radial basis function interpolation in trust-regions. SIAM Journal on Scientific Computing 30 (6), 3197–3219.

[64] Xiu, D., Hesthaven, J. S., 2005. High-order collocation methods for differential equations with random inputs. SIAM Journal on Scientific Computing 27 (3), 1118–1139.

[65] Xiu, D., Karniadakis, G. E., 2002. The Wiener–Askey polynomial chaos for stochastic differential equations. SIAM Journal on Scientific Computing 24 (2), 619–644.

[66] Yentsch, R., Gaitonde, D., August 1998. Exploratory simulations of the HIFiRE 2 scramjet flowpath. In: 48^{th} AIAA Joint Propulsion Conference. No. AIAA-2012-3772.